

# Using AI to Apply Computational Fluid Dynamics and Genetic Optimization to the Design of Airgun Moderators

A collaboration between Thomas Himes, Michael Erskine, and ChatGPT

Final Report

May 18, 2026



Ember Family moderator bodies tested in this project.

M. Erskine 2026 All Rights Reserved

# Authors and Collaboration

This paper is a collaboration between Thomas Himes, Michael Erskine, and ChatGPT. The human authors supplied the design problem, engineering judgment, test artifacts, model outputs, field data, and interpretive review. ChatGPT assisted with computational analysis, graphing, document drafting, and report assembly under human direction.

**Thomas Himes.** Air gunner with decades of experience and formal competition. Mechanical Engineer with 20 plus years design and application experience in high pressure compressible fluid flow and an additional 15 years design, testing, application and marketing of rechargeable batteries during the evolution to Lithium-ion technology.

**Michael Erskine.** Shooter and Computer Scientist with fifty years SIGINT collection, analysis, reporting and DoD/DARPA R&D experience.

Raw audio test data is available upon request at michael\_erskine@oldspooks.com.

## Executive Summary

The purpose of this project is to explore the subject expressed in the title: using AI to apply computational fluid dynamics concepts and genetic optimization to the design of airgun moderators. A study of the Ember Family was instituted to determine whether the genetic algorithm approach was viable and whether model-derived design trends could be connected to measured acoustic performance.

The project evaluates four related airgun moderator designs: Ember, Standard, Ideal, and Wraith. Each design was compared against a measured bare baseline using a calibrated microphone at approximately 116.5 ft from the muzzle. The primary measured comparison band is 300-3000 Hz because that band captures much of the principal audible report character and is less vulnerable to misleading emphasis from low-frequency rumble or high-frequency artifacts.

Measured results show that Wraith is the quietest device in the primary band, followed by Ideal, Standard, and Ember. A separate volume study shows that larger internal/acoustic volume correlates with improved measured suppression, but with diminishing incremental return as volume is increased.

The report does not claim mathematical proof that the true global optimum was discovered. It supports the more limited and defensible conclusion that the genetic algorithm found high-performing near-optimal designs under the implemented psycho-acoustic fitness model, and that measured field results support the viability of the approach.

Rank	Device	Measured bare-device difference 300-3000 Hz (dB)
1	Wraith	17.13
2	Ideal	15.17
3	Standard	13.93
4	Ember	11.50

## Background and Project Intent

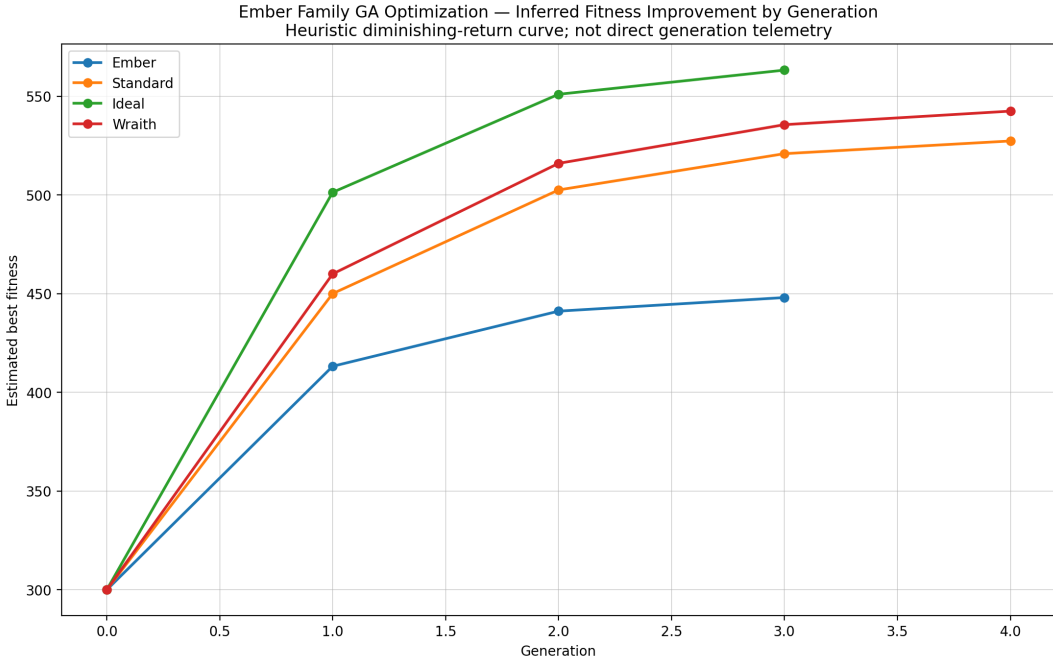
The Ember Family Moderators project connects three tasks: computational design, psycho-acoustic genetic optimization, and measured field validation. The genetic algorithm was used to search a large design space for moderator geometries that score well under a psycho-acoustic objective based on the perception of impulsive sound by mammals. The exact implementation details are not addressed in this public report; the report instead describes the general approach and evaluates whether the resulting designs produced measurable performance gains.

The genetic algorithm campaign used populations on the order of two million individuals. Standard and Wraith production runs evaluated 2.0 million candidates each, while Ember and Ideal used comparable 1.5 million candidate budgets. The fitness-improvement curves included here are heuristic diminishing-return reconstructions from the saved JSON summaries, not direct generation-by-generation telemetry. They support a practical near-convergence interpretation but do not prove discovery of a global optimum.

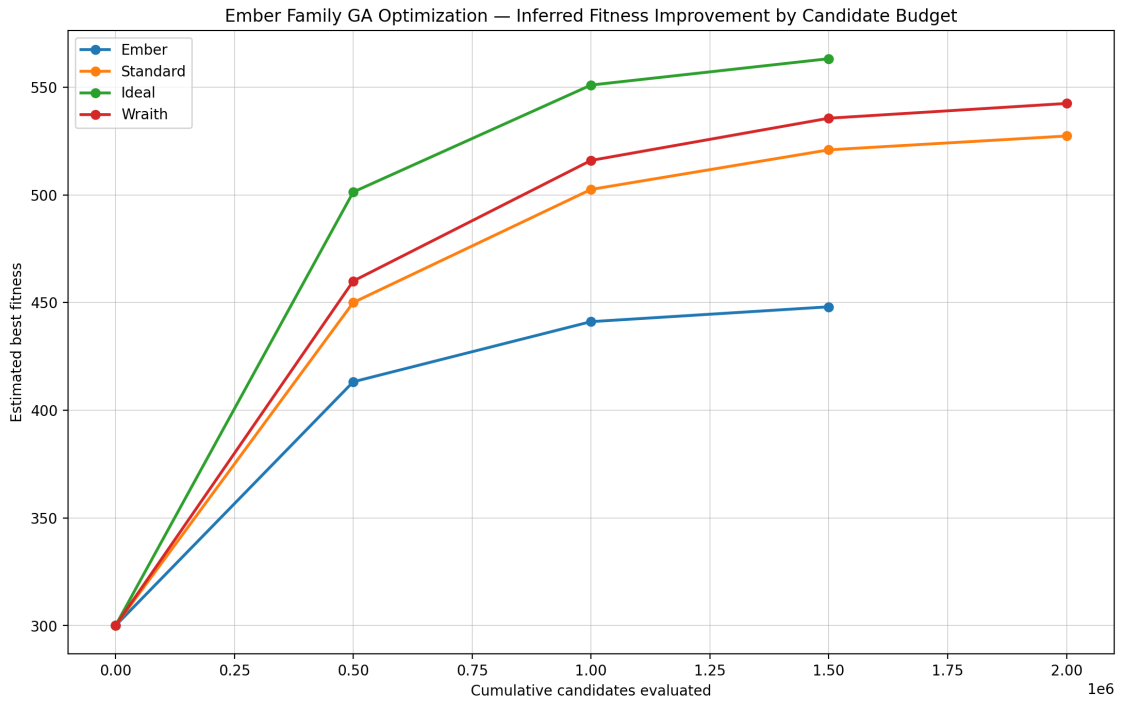
Device	Best fitness	Runs	Generations	Population	Total candidates
Ember	448.02	2	3	250,000	1,500,000
Standard	527.32	2	4	250,000	2,000,000

Device	Best fitness	Runs	Generations	Population	Total candidates
Ideal	563.21	2	3	250,000	1,500,000
Wraith	542.43	2	4	250,000	2,000,000

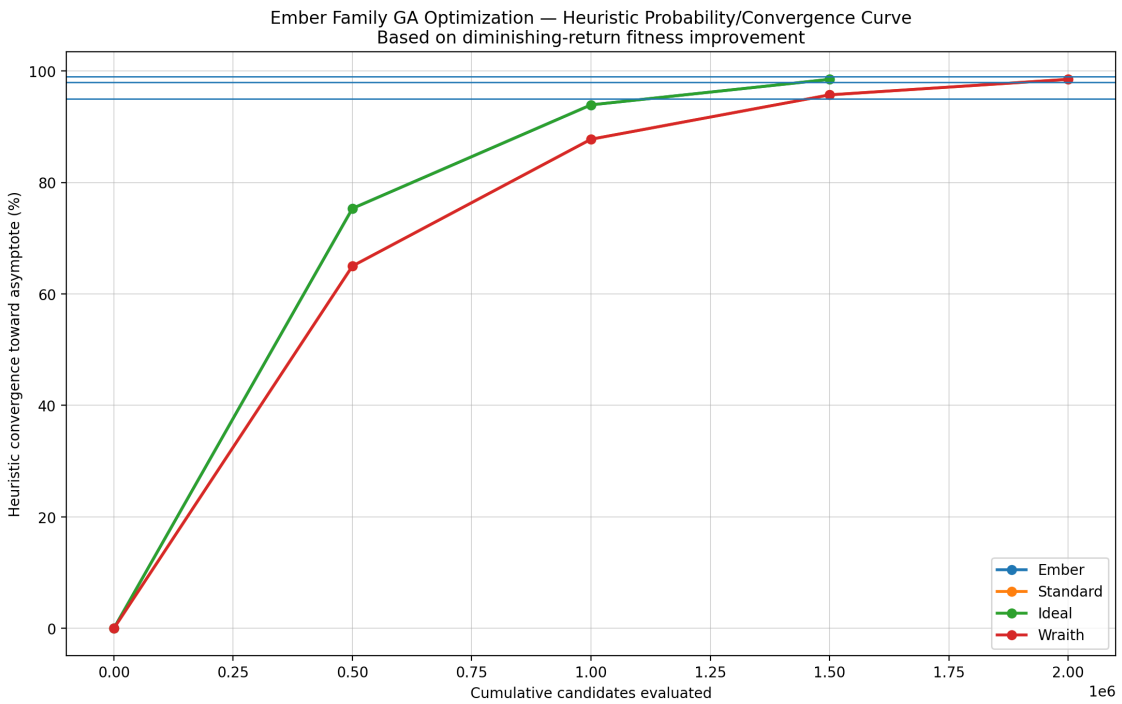
The inferred convergence heuristic treats the final best solution as approximately 98.5 percent converged toward a local asymptotic optimum. That value is an engineering confidence heuristic, not a mathematical probability that the global optimum was discovered.



Inferred GA fitness improvement by generation. Curve is heuristic; JSON files did not contain direct per-generation telemetry.



Inferred GA fitness improvement by cumulative candidate budget.



Heuristic convergence toward an inferred local asymptote.

# Accuracy Results

The target was shot at 42 yards with an Avenger air rifle in .25 caliber. The pellet was the JSB 25.39 grain pellet at 810 feet per second. Wind was light and variable, moving right to left across the range. Pellets were not sorted or inspected; they were selected directly from the tin in five-shot lots.

All shots on the target were fired consecutively. There were no sighters and no throw-away shots. Total shooting time for the target was less than one-half hour. The Wraith target in the lower right produced a flier due to a gust of wind; a sixth shot was then fired at that target only. All other groups are five shots each.



Accuracy target fired during the Ember Family Moderators test session.

# Moderators Tested



Ember Family moderator bodies used in the comparative test.

# Test Configuration and Data Scope

The report uses calibrated microphone recordings for measured signatures. Audio files were treated as mono. Although the WAV files were written as 24-bit / 192 kHz, the native/effective capture path was treated as 48 kHz and conclusions were limited to 20 Hz-20 kHz. The calibrated microphone was located approximately 38.83 yards, 116.5 ft, or 35.5 m from the source.

For this report, geometry is documented for traceability but no geometry-based correction is applied to the standard evaluation graphs. These are field measurements rather than anechoic laboratory measurements, so the strongest claims are relative comparisons among devices measured in the same study context.

## Analysis Procedure

The study-specific impulse window is -3 ms to +35 ms around the analysis center. Five impulses were used for each condition. Four suspect windows were adjusted after review: Standard shots 3 and 4, and Ideal shots 1 and 3. Adjustment moved only the analysis center of the suspect window so that the dominant peak aligned with the median dominant-peak offset of non-suspect shots for the same device.

Device	Shot	Original time	Adjusted center	Shift
Standard	3	75.209 s	75.200958 s	-8.04 ms
Standard	4	90.189 s	90.180958 s	-8.04 ms
Ideal	1	41.159 s	41.149958 s	-9.04 ms
Ideal	3	68.898 s	68.888625 s	-9.38 ms

A Standard Evaluation Graph is a corrected signature-view graph. It includes measured bare average signature, measured device average signature, and projected device signature derived from JSON TL data. The JSON TL is first shifted to align with measured actual suppression and is then converted into projected device signature by subtracting normalized TL from the measured bare signature.

## Device Design Summary and Volume


Device	Length mm	OD mm	Frustums	JSON volume mL	Envelope volume mL
Ember	106	33	3	85.07	90.66
Standard	135	35	5	122.73	129.89
Ideal	165	35	5	150.00	158.75
Wraith	172	40	5	207.02	216.14

## Measured Results

Device	20-20000	300-3000	1000-5000	2000-4000	8000-12000
Ember	7.65	11.50	12.49	14.24	-4.32
Standard	11.22	13.93	15.63	16.61	6.11
Ideal	12.59	15.17	18.01	19.86	4.86
Wraith	14.12	17.13	18.54	18.45	6.70

The adjusted-window results preserve the same ranking established before adjustment: Wraith is the strongest measured performer in the primary band, followed by Ideal, Standard, and Ember. Adjustment reduced the measured Standard and Ideal values but did not change the final order.

Ember Family Moderators — Adjusted Unnormalized Signature Comparison  
Black mic; adjusted suspect windows; window -3 ms to +35 ms




No per-device normalization. Suspect Standard/Ideal windows adjusted by dominant-peak alignment.

Adjusted unnormalized signature comparison for Bare, Ember, Standard, Ideal, and Wraith.

# Standard Evaluation Graphs

Ember: measured average bare-device difference over 300-3000 Hz is 11.50 dB. Psychoacoustic Noticeability Distance (PND) using the 300-3000 Hz RMS method and 8 dB short-impulse noticeability penalty is 104 yd.

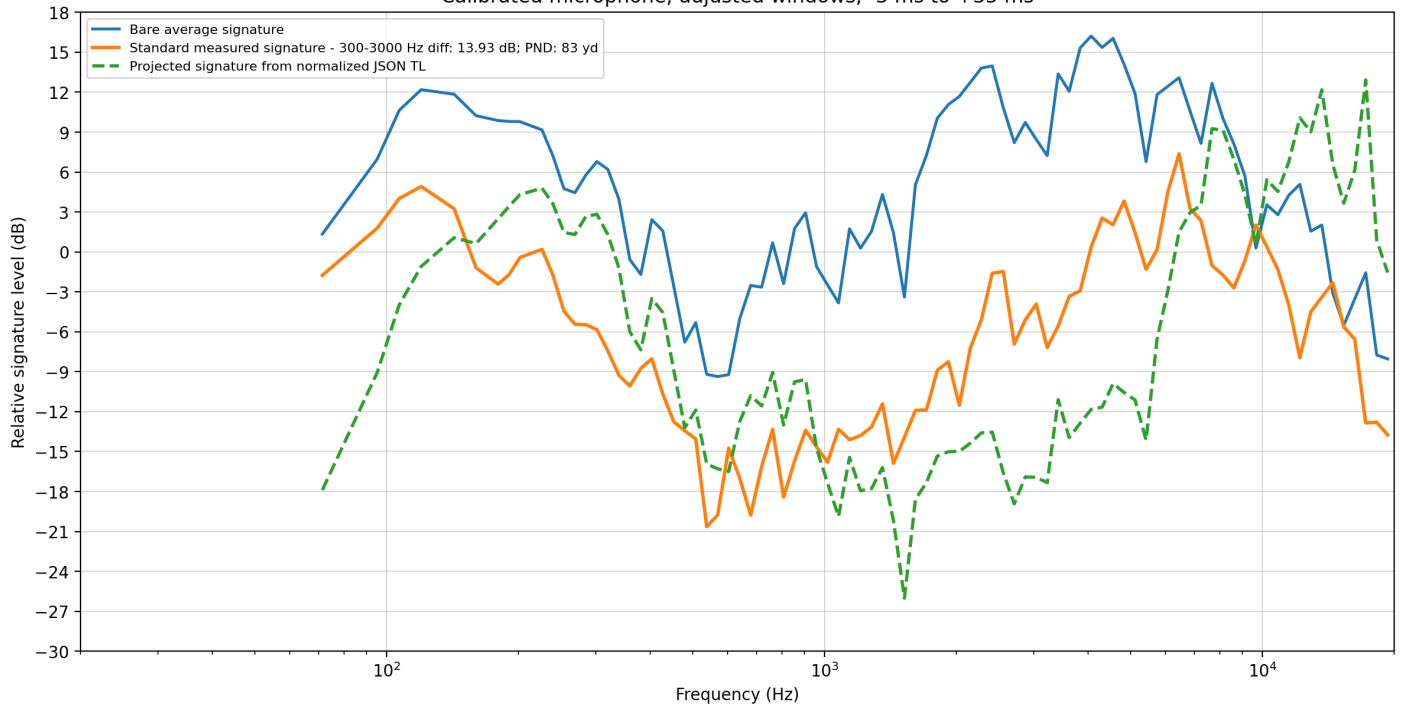


PND = RMS shot level above common 300-3000 Hz ambient RMS minus 8 dB short-impulse noticeability penalty; distance from 116.5 ft using  $20\log_{10}$  spreading.

Ember standard evaluation with measured bare signature, measured device signature, projected signature, and PND note.

Standard: measured average bare-device difference over 300-3000 Hz is 13.93 dB. Psychoacoustic Noticeability Distance (PND) using the 300-3000 Hz RMS method and 8 dB short-impulse noticeability penalty is 83 yd.


Standard Standard Evaluation  
 Calibrated microphone; adjusted windows; -3 ms to +35 ms



PND = RMS shot level above common 300-3000 Hz ambient RMS minus 8 dB short-impulse noticeability penalty; distance from 116.5 ft using 20log10 spreading.

Standard standard evaluation with measured bare signature, measured device signature, projected signature, and PND note.

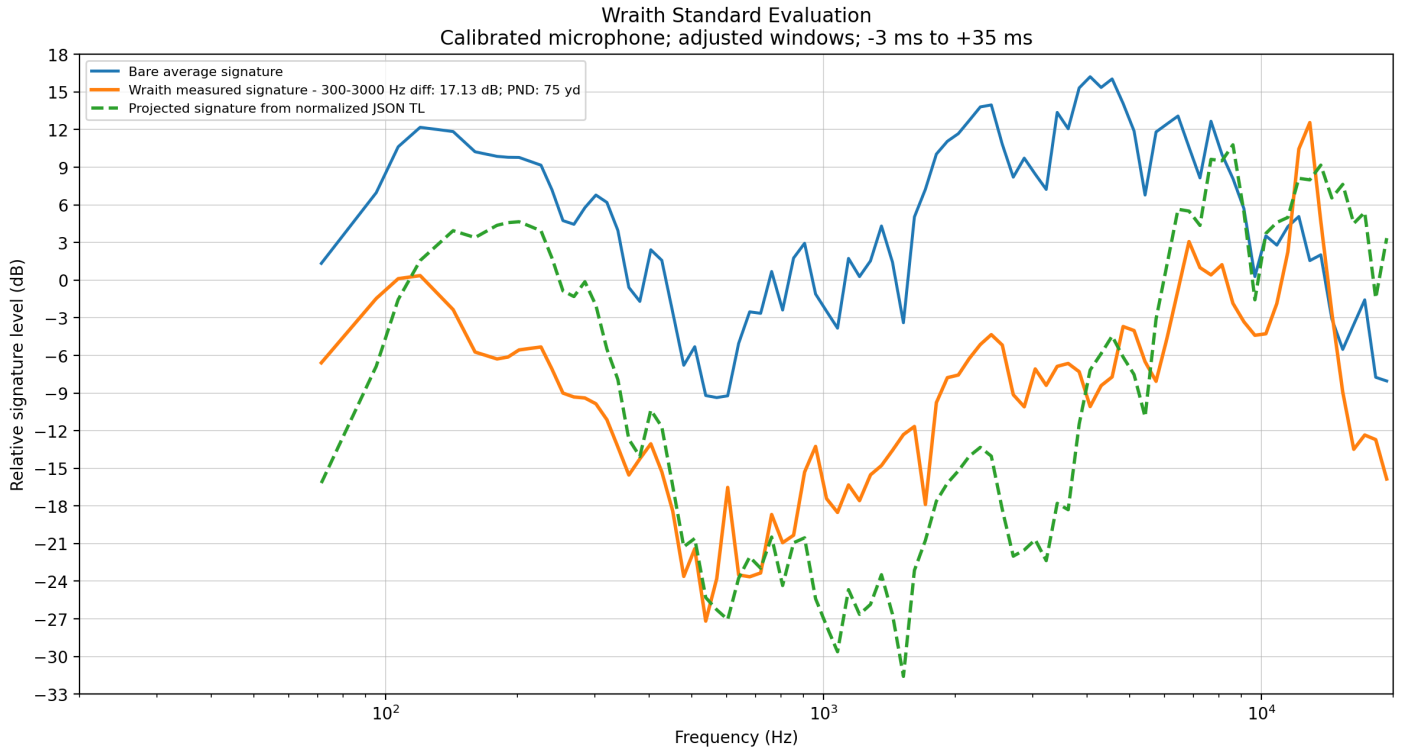
Ideal: measured average bare-device difference over 300-3000 Hz is 15.17 dB. Psychoacoustic Noticeability Distance (PND) using the 300-3000 Hz RMS method and 8 dB short-impulse noticeability penalty is 71 yd.



PND = RMS shot level above common 300-3000 Hz ambient RMS minus 8 dB short-impulse noticeability penalty; distance from 116.5 ft using 20log10 spreading.

Ideal standard evaluation with measured bare signature, measured device signature, projected signature, and PND note.

Wraith: measured average bare-device difference over 300-3000 Hz is 17.13 dB. Psychoacoustic Noticeability Distance (PND) using the 300-3000 Hz RMS method and 8 dB short-impulse noticeability penalty is 75 yd.



PND = RMS shot level above common 300-3000 Hz ambient RMS minus 8 dB short-impulse noticeability penalty; distance from 116.5 ft using 20log10 spreading.

## Projection vs Measurement

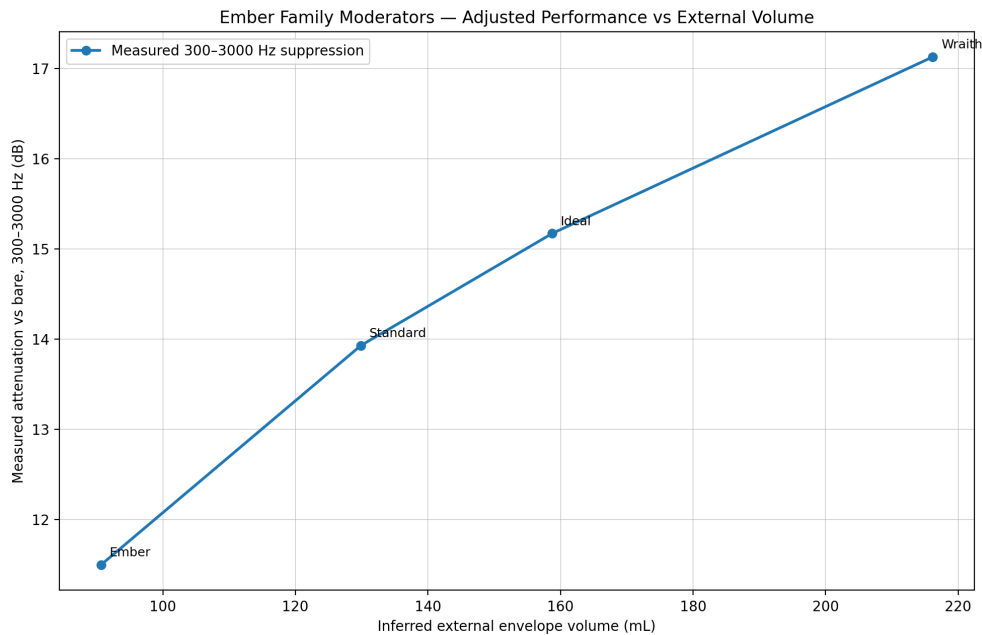
Device	Projection offset dB	Projected signature RMS error dB
Ember	-22.97	6.09
Standard	-24.88	8.85
Ideal	-29.02	7.96
Wraith	-28.61	8.60

The required offsets are large and negative, so raw JSON TL magnitude is not reported as measured performance. After conversion, the projected signatures are useful for shape comparison and model interpretation. Measured signatures remain the controlling evidence for quietness ranking.

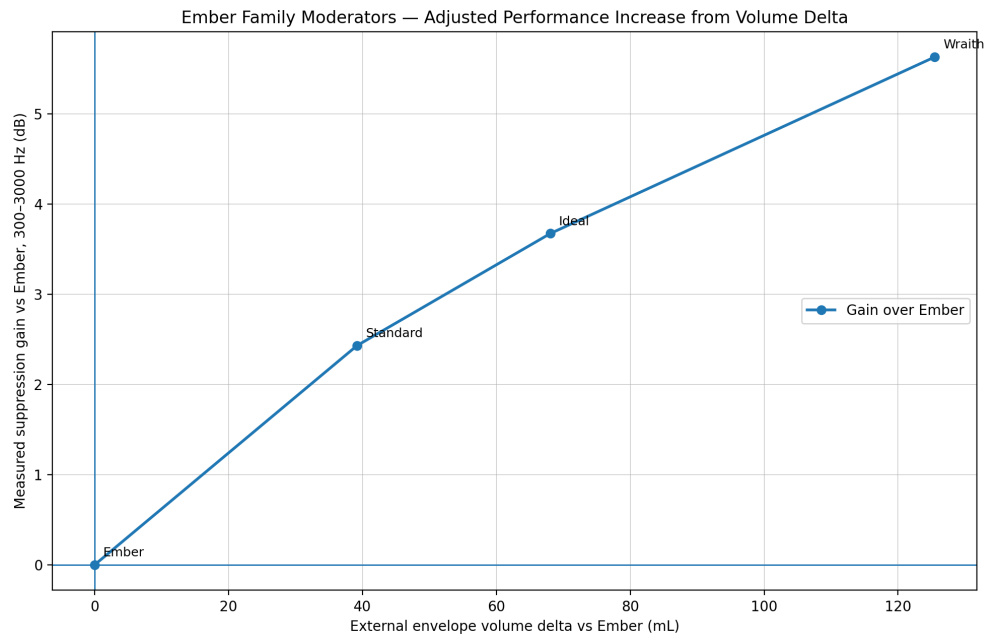
## Volume Delta and Performance Increase

Device	Ext. volume mL	Internal volume mL	Supp. 300-3000	Gain vs Ember	Gain/100 mL ext.
Ember	90.66	85.07	11.50	0.00	
Standard	129.89	122.73	13.93	2.43	6.20
Ideal	158.75	150.00	15.17	3.67	5.40
Wraith	216.14	207.02	17.13	5.63	4.49

The volume study shows increasing measured suppression with increasing internal/acoustic volume. The incremental dB gain per added 100 mL decreases as the devices become larger, supporting a diminishing-returns interpretation.



Adjusted performance versus inferred external envelope volume.



Adjusted performance increase versus volume delta from Ember.

# Psychoacoustic Noticeability Metric

This report introduces a field-useful metric for moderator and silencer evaluation: the Psychoacoustic Noticeability Distance, abbreviated PND. The purpose of this metric is to estimate the distance at which a short-duration moderated report becomes difficult for an unaided human listener to notice against the ambient acoustic background. It is not a laboratory SPL certification and not a guarantee of detection or non-detection. It is a repeatable comparative metric intended to improve practical silencer evaluation by accounting for frequency band, impulse duration, and signal power above ambient noise.

The PND calculation is designed to make the signal and ambient quantities apples-to-apples. Peak values are not used. The method uses 300-3000 Hz band-limited RMS for both the shot signal and the ambient background. This places both quantities in the same frequency band, the same time-domain RMS framework, and the same recording-domain dB scale.

Traditional peak-level comparisons can exaggerate practical detectability because they compare a very brief impulse peak against a longer-duration ambient noise condition. Human hearing does not respond only to instantaneous peak pressure. Detection of a short impulse depends on energy accumulation over time, masking by background noise, and frequency content. For this reason, the PND metric uses RMS energy in the principal audible report band rather than peak amplitude.

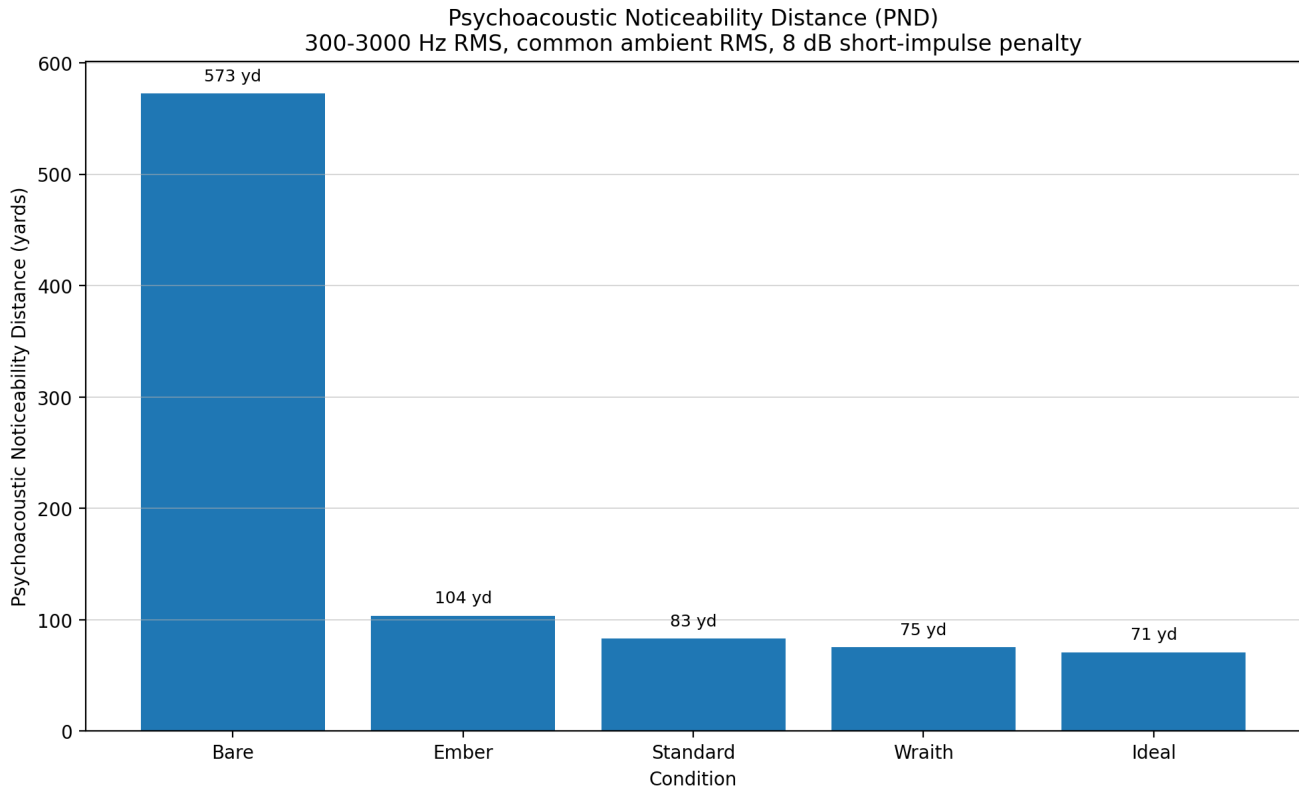
For this study the metric is defined in the 300-3000 Hz band using the 38 ms impulse window. The common ambient reference is the average 300-3000 Hz RMS ambient level measured from non-shot sections of the test recordings. The common 300-3000 Hz ambient RMS value is -62.50 dB on the recording-domain scale.

Because the event is a short impulse, an 8 dB short-impulse noticeability penalty is applied. This approximates reduced detectability of a 38 ms impulse relative to the roughly 200-300 ms temporal integration interval commonly associated with human auditory detection and loudness perception.

Psychoacoustic Noticeability Margin (PNM) = shot\_RMS\_300-3000Hz\_dB - ambient\_RMS\_300-3000Hz\_dB - 8 dB.

Psychoacoustic Noticeability Distance (PND) = measured\_distance x 10<sup>(PNM / 20)</sup>.

Device	Shot above ambient dB	PNM dB	PND ft	PND yd	PND miles
Bare	31.37	23.37	1718	573	0.33
Ember	16.52	8.52	311	104	0.06
Standard	14.60	6.60	249	83	0.05
Ideal	13.24	5.24	213	71	0.04
Wraith	13.75	5.75	226	75	0.04



Psychoacoustic Noticeability Distance using 300-3000 Hz RMS, common ambient RMS, and an 8 dB short-impulse noticeability penalty.

The PND value should be interpreted as a relative field estimate of the distance at which the report falls into the ambient acoustic background for an unaided listener. Real-world noticeability also depends on terrain, wind, foliage, reflections, listener orientation, expectation, hearing acuity, and the spectral character of the ambient noise. PND is intended to move silencer evaluation away from simple peak reduction and toward the more practical question: how far away does the moderated report remain noticeable above the actual ambient sound field?

## Validity and Supportability

The results are supportable as a relative field comparison because all devices were compared against the same bare baseline, five impulses were used per condition, the same analysis window was used, the same calibrated microphone and audio assumptions were applied, and no per-device normalization was applied to measured signatures for quietness ranking.

Some clipping occurred in the bare baseline samples recorded at approximately 116 ft from the muzzle. No clipping was observed in any of the tested moderator samples. The test is therefore deemed valid for the comparative purpose. Bare-only clipping can reduce or distort apparent bare energy, so the measured moderator results are at least as good as stated rather than overstated by moderator clipping.

## Limitations

This is a field comparison, not an anechoic laboratory test. Absolute SPL certification is not claimed. The files are written at 192 kHz but are interpreted as effective 48 kHz capture. Geometry is documented for context but not used as a correction term in the final standard evaluation. Model projection data is secondary and is interpreted only after TL-to-signature conversion.

## Conclusions

The project demonstrates that an AI-assisted, psycho-acoustic genetic optimization approach can produce airgun moderator designs with measured field performance differences that follow expected design trends. In the Ember Family study, Wraith measured quietest in the primary 300-3000 Hz band, followed by Ideal, Standard, and Ember. Volume increases generally improved measured suppression but with diminishing returns.

The genetic algorithm results, volume study, standard evaluation graphs, and Psychoacoustic Noticeability Distance analysis together support the viability of the approach. The report does not claim that the true global optimum was mathematically proven. It supports a practical engineering conclusion: the GA approach found high-performing candidate designs, and the measured data validates the approach as useful for further moderator design work.

# Appendix A - Source Files and Generated Artifacts

Artifact	File
Psychoacoustic Noticeability final summary	ember_family_pnd_final_common_ambient_summary.csv
Final Ember standard evaluation	final_ember_standard_eval_pnd.png
Final Standard standard evaluation	final_standard_standard_eval_pnd.png
Final Ideal standard evaluation	final_ideal_standard_eval_pnd.png
Final Wraith standard evaluation	final_wraith_standard_eval_pnd.png
Adjusted combined signature graph	ember_family_adjusted_unnormalized_signature_comparison_all.png
Adjusted volume/performance table	ember_family_adjusted_volume_performance_table.csv
Psychoacoustic Noticeability graph	ember_family_psychoacoustic_notice_distance.png

Source data consists of user-provided WAV files, user-provided target and moderator photographs, and JSON model outputs generated by the optimization process.